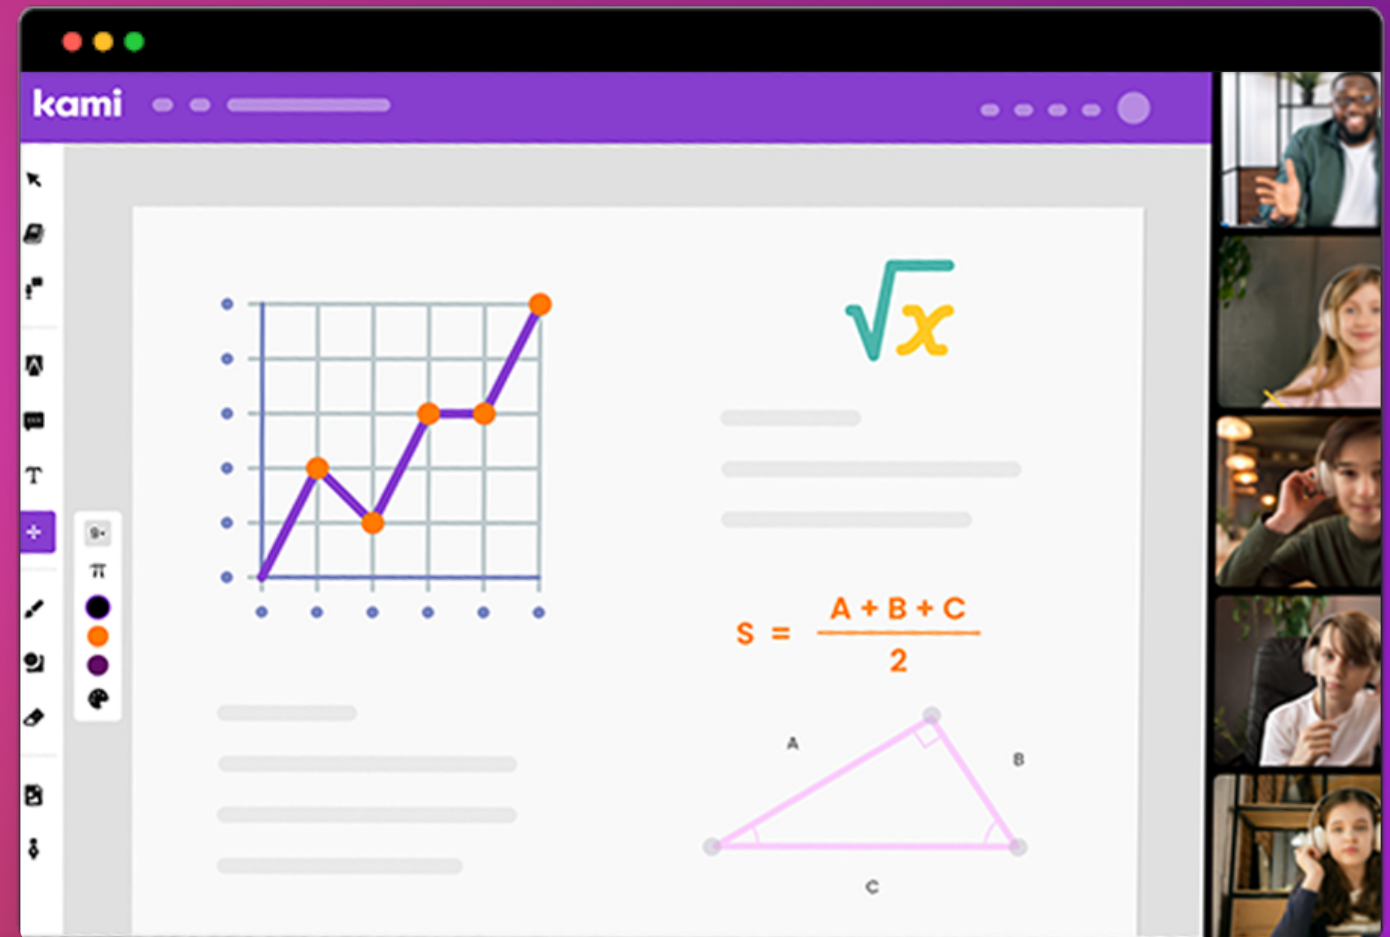


[CASE STUDY]

kami

A lesson in scaling:
educating millions of
students during a pandemic



The screenshot displays the Kami online learning interface. The main content area shows a math lesson with several elements:

- A line graph on the left with a purple line and orange data points on a grid.
- A square root symbol \sqrt{x} in the top right.
- A formula $S = \frac{A+B+C}{2}$ in the middle right.
- A right-angled triangle with vertices labeled A, B, and C at the bottom right.

The interface includes a toolbar on the left with various drawing and editing tools, and a vertical video feed on the right showing multiple participants in a virtual classroom.



How Kami migrated their mission-critical workload from Postgres to CockroachDB in 4 months

INDUSTRY :

Education

CHALLENGES :

Scale to accommodate over 1M users per week.

SOLUTION :

Migrated from Postgres to CockroachDB to grow biz 20x.

About Kami

Kami was founded in 2013 when the four New Zealand-based co-founders saw a need for a more collaborative way to take notes and avoid the need to print, which was costing around \$8B a year worldwide. Kami is now the world's #1 digital classroom application; it creates a flexible and collaborative learning environment for millions of teachers and students across the globe.

The education industry is slow when it comes to tech adoption despite the fact that they spend around \$700B a year on software purchases. In fact, 60-70% of software goes unused in schools. In order to help close this gap, Kami's platform allows teachers to digitally display documents on the student's laptops or Chromebooks instead of printing individual handouts. Teachers and students can also take notes directly on the platform.

Even before the pandemic, Kami's business was doing well and there was clearly a need for this type of digital alternative to printing in schools. But in 2020 when the pandemic hit, schools closed and students were forced to learn from home. The Kami team was glad to see their platform was able to keep students and teachers connected... until they saw what it meant to scale their Postgres database in the face of exponential growth.

At the peak of the pandemic, Kami was growing at a rate of over 1M users per week, and their Postgres database simply couldn't keep up. They doubled their database size only to find that it bought them two weeks of runway. They spent sleepless nights removing performance bottlenecks that earned them mere days of uptime.

As the team put out fire after fire, the true weight of the situation became clear: failing to scale didn't just mean going out of business, it meant kids around the world would fall behind.

Fortunately, their team found a solution: CockroachDB. After a live migration, they were thrilled to see that CockroachDB seamlessly helped them grow 20x in six months. More importantly, Kami was able to help teachers and kids around the world (even making their platform free!) while their competitors were forced to shut their doors.



Forecasting growth & a storm ahead

Kami has a lean engineering team of seven members, and they originally built their platform on Postgres. That served them well for many years, and they were able to scale from zero users to several million. However, when businesses like Kami start to see tremendous growth, Postgres begins to hit limitations.

The team made a few attempts to scale and optimize the database manually. They started with upsizing their PostgreSQL primary to 160vcpu, 4TB RAM, which was the biggest server they could get. However, the write-heavy data from user-driven content was becoming too much to handle. They scaled up write capacity, but they couldn't scale horizontally anymore. There were also operational issues with implementing transaction ID wraparounds. All the written data needed to be vacuumed after 2 billion transactions. But they were pushing close to a billion transactions every day!

For their setup, connection pooling was necessary, but they were hitting their scaling limits here as well. With over 150 active connections from pooler to PG, performance would decrease even with 160 vCPUs on the Postgres server and plenty of spare CPU.

Ultimately they could get up to 120,000 queries per second through Postgres and support ~4 million active users. That was not going to be enough for the new school year starting in ten weeks! They could optimize it even further, but at best that would be a 2x or 5x gain, and after that Kami would be stuck.

The team felt like they hit their absolute limit with Postgres and needed a new solution, FAST. Given Postgres' limitations, when it came to choosing a new database, the Kami team had the following requirements:

- ✓ Ability to scale horizontally and add capacity over time as needed
- ✓ Accommodate write-heavy workloads
- ✓ Compatible with Postgres so the transition would be fast
- ✓ Simple to operate, especially simple to scale
- ✓ A winning technology that would grow with them 10+ years

The team considered a number of different databases, but they didn't have time to run performance tests. They looked at Yugabyte but felt there were a few aspects where they lacked in maturity (especially schema changes). They also looked at Citus which is a Postgres extension designed to help you scale out. However, Citus had some limitations around the write traffic (using a single node to coordinate writes) and could not scale horizontally for writes.

Luckily, someone on the team remembered reading about a distributed SQL database named CockroachDB on Hacker News a few years ago, and decided to investigate further.

30 million

USERS

20x

GROWTH

190+

COUNTRIES

~1 billion

TRANSACTIONS/DAY



Trusting a famously resilient creature: CockroachDB

Usage and sales were doubling month over month and the database was facing immense pressure already. Co-founder and Chief Technology Officer Jordan Thoms was ready to start evaluating CockroachDB when fellow Co-founder and Chief Revenue Officer Bob Drummond reminded him that he needed a solution fast. He had early insight into sales, and knew that Kami was expecting 10x load growth within the next two months.

Jordan reached out to Cockroach Labs in May 2020 and kicked off a brief evaluation. The Kami team liked that they could continue to use SQL and that CockroachDB offers the ability to partition data by location so they can easily cater to a global audience and scale their environment across regions. Additionally, they were impressed with CockroachDB's performance when it came to managing a dataset that is continuously being updated.

Ultimately the biggest advantage of CockroachDB was its ability to easily scale and provide a resilient foundation for Kami's platform. They were on a mission to help any teacher and student around the world use the Kami platform—that meant that they needed it to always be on and available.


Even though the decision was made quickly, the Kami team felt like CockroachDB was the right solution to scale with them over the next 10+ years. Now it was time to migrate their most critical workloads from Postgres to CockroachDB.

“ ”

“We knew how complicated it would have been to stay with Postgres and set up sharding. There would have been a constant drag of managing multiple shards of data. No one on our team wanted to go through that manual labor. And even if we did set up sharding, we weren't going to be able to grow our business 10x.”

– **Jordan Thoms**
co-founder, CTO





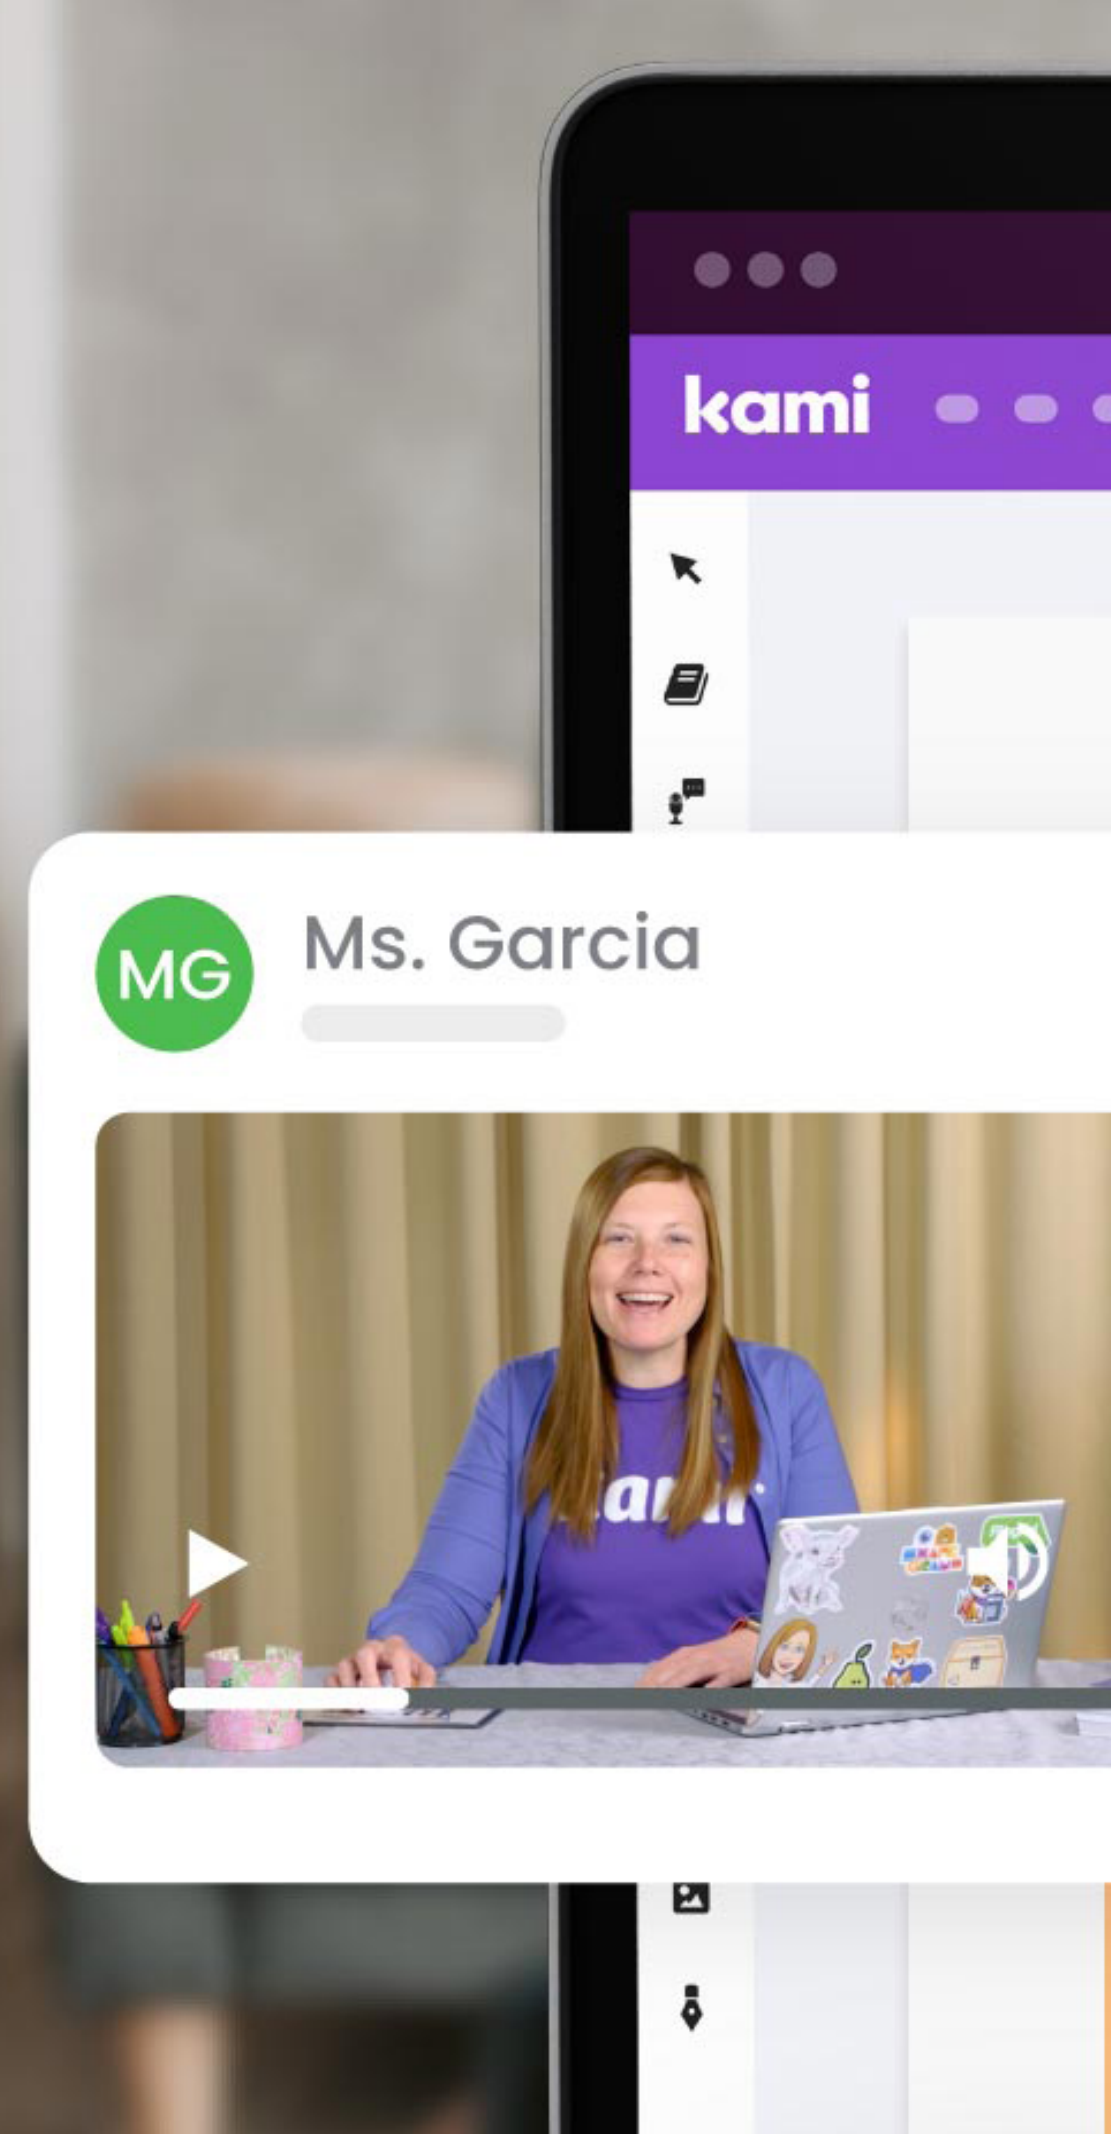
Executing a live migration

A “traditional” database migration strategy would include setting up a production system to write to the new database and the old database, and read only from the old one. In the background, you migrate previous data from the old database to the new one, and ultimately switch reads over to the new database and turn writes off to the old database. If you have the time, it’s smart to limit the effects of a migration to a small set of users at first before gradually rolling out to a larger audience.

But this was not a traditional situation, and Kami simply did not have the time.

Instead they needed to do a live migration which meant setting up a load and correctness test. Initially all documents were written to Postgres, but they progressively enabled ‘New DB’ (i.e. CockroachDB) for more users over time. Then they started to mirror all writes to the new database (CockroachDB) and old database (Postgres), and they would read from both databases and check that the result matched. Mismatches were collected into a new table and investigated. This approach allowed them to gradually gain confidence in the performance and correctness of CockroachDB with real-world usage.

The migration started in early June and concluded in mid-September, right in time for the beginning of the 2020 school year. The team was under an incredible amount of pressure, but they were successful!

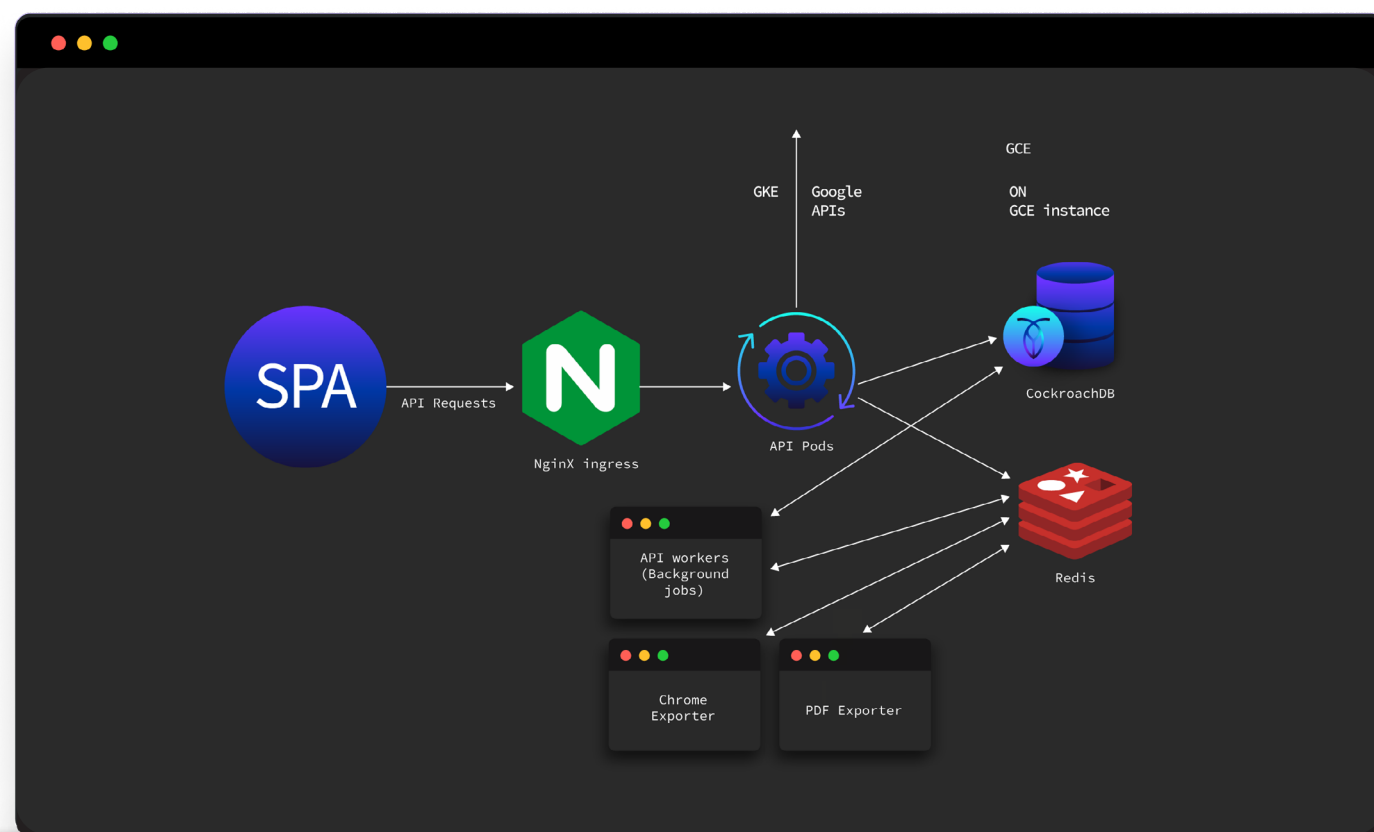


A scalable infrastructure for the future

Fast forward to summer 2021 and their infrastructure is built to sustain over 200K annotations per minute. All of this annotation (and document) data is being stored in CockroachDB. This is a mission-critical, write-heavy workload. Eventually, Kami will move user login data to CockroachDB as well.

Today, Kami's primary API sees up to 25K requests per second, generating over 90K qps to CockroachDB. requests per second. Their Google Kubernetes Engine environment auto-scales daily up to 3,000 cores and 11TB RAM. They are still using Postgres to some extent, with plans to fully migrate over time. However, they've already moved the most valuable tables to CockroachDB and are running a 50 node cluster with 16vCPUs each.

Here's an overview of their setup:



“ ”

“After switching to CockroachDB, our team is much more productive and we didn't have to recruit a bunch of Postgres DBAs. We are really pleased with its performance and our ability to massively scale. It's the winning technology for the future.”

– **Jordan Thoms**
co-founder, CTO



What's next for Kami

Because of the pandemic, 5+ years of digital education trends were condensed into about seven months (March - September 2020). Fortunately, Kami was able to complete a fast migration to a powerful new database that made their platform available to over 27 million users across the globe.

Kami's business grew 20x in the six months after the 2020-2021 school year began. Right now, usage is heaviest in the U.S. (which accounts for over half of the schools using the platform), but Kami is seeing substantial growth in Canada, UK, other European countries, and South Asia.

The Kami team is glad to see the education industry adopting new technologies and hopes it will positively impact teachers and students' lives well into the future.

You can learn more about Kami at their website: www.kamiapp.com.

Ready to get started?

Go hands-on with 100% free CockroachDB Serverless. Spin up your first cluster in just a few clicks.

